arXiv:1703.10584v1 [cs.CV] 30 Mar 2017

Geometric Affordances from a Single Example via the Interaction Tensor

Eduardo Ruiz, Walterio Mayol-Cuevas
Department of Computer Science
University of Bristol, UK
{er13827,wmayol}@bristol.ac.uk

Abstract

This paper develops and evaluates a new tensor field representation to express the geometric affordance of one object over another. We expand the well known bisector surface representation to one that is weight-driven and that retains the provenance of surface points with directional vectors. We also incorporate the notion of affordance keypoints which allow for faster decisions at a point of query and with a compact and straightforward descriptor. Using a single interaction example, we are able to generalize to previously-unseen scenarios; both synthetic and also real scenes captured with RGBD sensors. We show how our interaction tensor allows for significantly better performance over alternative formulations. Evaluations also include crowdsourcing comparisons that confirm the validity of our affordance proposals, which agree on average 84% of the time with human judgments, and which is 20-40% better than the baseline methods.

1. Introduction

Perhaps the most fundamental question about Vision is *what is it for?* From the early propositions to computationally address this by D. Marr [15], the path has often been assumed to aim to, or at least require to, recover geometric information about the environment. This has directed much effort towards a particular spatial representation of the world, but at the same time, has moved effort away from the rationale that drove Vision to its high level of utility.

The view posed by J.J. Gibson however [7], calls for a visual perception that is there to help the perceiving agent to interact with the world. Specifically, through the coining of the term *affordance*, visual perception is described as a process to understand what can be done where. Such a representation of the world is immediately useful as by definition it is one that already takes into account what the agent is capable of.

ER is supported by the Mexican Council for Science and Technology (CONACYT) scholarship scheme.

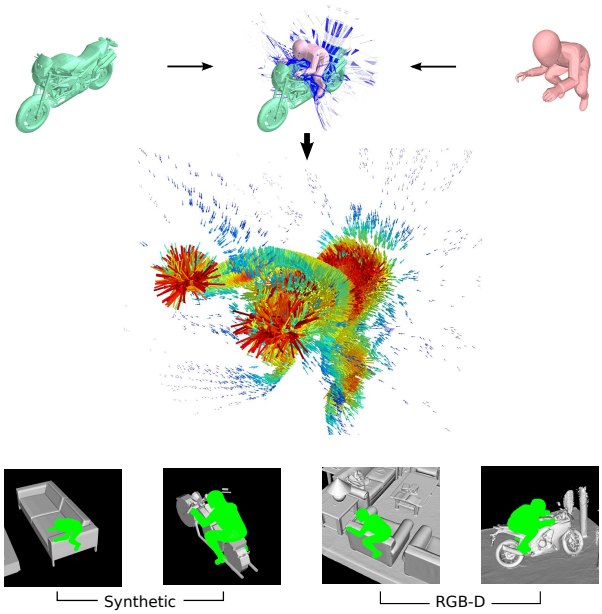

Figure 1: Our affordance tensor (centre) describes the interaction pair of "riding" (top row). This allows us to predict affordance locations on previously unseen scenes (bottom row), even under changes of geometry and from a single example. The geometric affordance we estimate agrees with judgments of mechanical turk markers, and is able to answer something like: *where would the kids pretend to ride a motorbike on the living room?*

Gibson also argued that affordances are "immediate" to perceive. This has often been misread as a call to ignore the relevance of the representation [28]. But direct perception of affordances does not mean that no intermediate processing should take place, including perhaps requiring a 3D reconstruction; for instance, it has been shown that the dorsal stream in the visual cortex of the brain, computes edge detection, depth, 3D surface and axis representations [19]. Importantly, we argue that, the direct nature of affordance perception rather motivates methods that are able to immediately transfer what has been learned to other objects and

places after a small number or even a single observation of the affordance.

Being able to determine affordances can have profound implications for visual systems. It can in principle liberate the computational approach to visual processing from the focus on objects and their arbitrary labels which have to be extensively learned. To learn an affordance is not to classify an object [7], since a cup is not only for drinking but also a paperweight, or even a tool to build sand castles.

While the concept of an affordance can appear elusive to express, an affordance is necessarily the result of the composition between the world and the agent. And one that is ultimately designed to be useful for the perceiving agent. Understanding and modeling this interaction between agent and the world is of central focus to our work and the one that we aim at.

We here concentrate on the subclass of affordances between rigid objects. Affordances such as where can I hang this?, place this, ride, fill, and similar. We do this by specifying a geometry-driven interaction tensor that aims to capture the way in which the affordance manifests between a pair of objects.

Importantly, using only a single example, we then detect other viable places for such affordances in previously unseen scenes. We evaluate with both synthetic and real scenes.

Our approach is inspired by the well established concept of bisector surfaces (see [21] for an introduction), and their recent exploits for scene indexing [29]. Here we extend these concepts by directly enhancing the bisector surface points with rational weighting, provenance vectors and the concept of affordance keypoints. All which results in a richer vector field or tensor. Our contributions in this paper can be outlined as follows:

- We extend the notion of the bisector surface to a weighted vector field —an interaction tensor field.
- Show how this tensor with direct, sparse sampling, allows for the determination of geometrically similar interactions even from a single example, and is better than existing formulations.
- Introduce the notion of *affordance keypoints* which serve to more quickly judge the likelihood of an affordance at a query point.
- Evaluate with both synthetic and real scenes from RGBD mapped areas.
- We validate results with crowdsourced judgments.

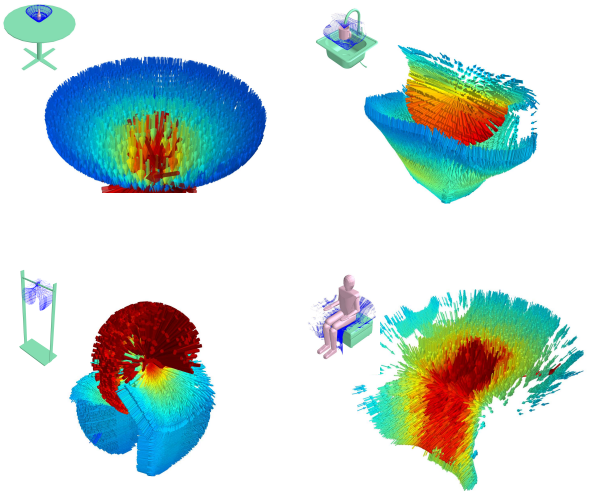

Figure 2: Interaction tensor examples of 4 affordances. Starting from the top-left in clockwise direction: *placing* a bottle, *filling* a mug, *sitting* and *hanging* a coat hanger.

2. Related work

Affordance detection has been studied in recent years in computer vision and robotics. Briefly speaking, affordance knowledge has been incorporated in learning systems that use data from demonstrations of interaction, robot self exploration and static labeled imagery. In terms of the applications, the approaches include semantic scene understanding, grasp learning, gesture recognition, object segmentation and planning in goal-directed tasks.

An important body of research comes from the developmental robotics field [16]. The core of these approaches is the representation and learning of actions and predicting the consequences of these over a set of objects. These approaches use visual features describing shape, color, size and relative distances to capture object properties and effects. Using robot self-exploration and human demonstrations the systems benefit from single-object affordances to execute more complex interactions and execute a plan (task planning). For instance, [14] shows a robot learning in a self-supervised manner to use a tool by observing the effects of its actions on other objects.

Another line of research that has benefited from affordance learning is Human-Robot Interaction [26, 20, 22, 12, 27, 3, 10]. In these studies the main goal is to to perform action recognition in a robot observing humans, usually to predict or anticipate human activities, and in this way assist humans better while they perform everyday tasks.

Work has also been done using static imagery, where the affordance or interaction is provided as a label rather than demonstrated. [6, 27, 31, 4] based their work on labeled 2D imagery to predict functional regions or attributes on every day objects.

A body of research closer to our approach is the one exploiting 3D information to learn and predict affordances of objects in the environment. In [1], the concept of 0-order affordance is introduced to refer *hidden* affordances that can be found on an object but not in its current pose. Amongst the affordances studied are rollable, containment, liquid-containment, unstable, stackable-onto and sittable. In [9] a physics-based simulation on CAD models of objects is used to learn three functional classes: drinking vessel, table and sittable. Using geometric features on RGB-D data [11] presents a segmentation algorithm that learns and predicts affordances such as pushable, liftable and graspable on indoor scenes. In [17, 18] RGB-D images are used to learn and predict functional regions such as grasp, contain, support and cut on objects placed on a table-top. Using RGB-D images of indoor scenes [24] perform segmentation for human actions such as walkable, sittable, lyable. Similarly, affordances are studied in [23, 8, 10] to map locations suitable for sitting, or laying down; particularly in these cases using human skeleton *hallucinated* on the different indoor scenes. Crucially, these previous methods are heavy in terms of requiring multiple learning examples, impose a particular parameterization such as detection of planes or shapes and or are highly specific to an object *e.g*. humanoid shapes. Our approach aims to address various of these limitations, namely relying on pre-parameterization of scene or objects and relying in numerous examples.

3. Our approach

As noted before, geometrical properties are very important in the modeling or representation of affordances. And this is validated by success of some of the related methods.

An interesting additional related work is [29], where an algorithm for 3D scene indexing is developed to capture hierarchical relationships among objects using Betti numbers. It proposes Interaction Bisector Surface curvature descriptors that are learned from multiple examples. There, it is shown the discriminative power of the Bisector Surface (BS) to characterize the relationships between sets of objects. The bisector surface B of two objects O_1, O_2 is the locus of points equidistant to the objects' surface. The bisector surface is an approximation of the Voronoi diagram for objects in a scene.

We extend the robustness of the BS by preserving information regarding the expected locations or areas in the 3D space that enable the interaction. This is what we call the Interaction Tensor (iT).

Furthermore, we are able to identify areas of high importance for each affordance based on the geometric and spatial relationships between the interacting objects. Briefly speaking, our method consists on computing the iT descriptor between a pair of objects of whose affordance we are investigating. Examples of the iT between pairs of objects are shown in Fig. 2.

As first step, the objects are placed simulating the interaction that they would have on real circumstances (affordance example); we then compute the BS produced by these two objects and preserve the **provenance vectors**. That is, the vector(s) that contributed to the computation of a given point in the BS. But note that the provenance vectors should not be confused with surface normal vectors on the BS, since the latter do not provide information regarding from where the BS points actually come from. This process generates the Interaction Tensor for the affordance simulated by the two interacting objects. At test time, we are able to predict affordance location candidates by approximating the iT on a previously unseen input scene. The method allows us to use a model of say a humanoid *skeleton* and predict *human affordances* such as sitting; similarly to [8, 10, 23, 24]. But importantly, it also allows us to build these tensors more generally for any other pair of objects.

Formally, given a bisector surface B formed by points b and an object O in the scene formed by points $o_i, ..., o_n$, the tensor field characterizing the interaction is defined as

$$\mathbf{iT}(B) = P\hat{i} + Q\hat{j} + R\hat{k} \tag{1}$$

where

$$P = \hat{G}(B)_{\hat{i}} - B_{\hat{i}}$$

$$Q = \hat{G}(B)_{\hat{j}} - B_{\hat{j}}$$

$$R = \hat{G}(B)_{\hat{k}} - B_{\hat{k}}$$

with

$$\hat{G}(B) = \underset{i}{\operatorname{argmin}} \|o_i - B\|_2$$

Additionally, we take into account the measurement of how important every location is in the iT. This is expressed as a weight as we will discuss later.

3.1. Computing the Interaction Tensor

The bisector surface between a pair of objects is computed similarly to [29]. Using 3D or CAD models of the interacting objects, the first step is to create dense point clouds by uniformly sampling points on the surfaces of the models. Then, the Voronoi diagram is computed for all these points; which produces a simplicial complex where polygon ridges are equidistant to the points that produced them. The Bisector Surface is comprised of ridges originated by points from different objects. In our experiments, we refer to the two interacting objects as **query-object** and **scene-object** (or scene) respectively. The query-object is the one with a known affordance; a mug which affords *filling*, is an example of an affordance query in our setup. A scene-object is the second part of the interaction; this could be a second object or part of a scene or furniture that allows the affordance to take place. Using the same mug filling example,

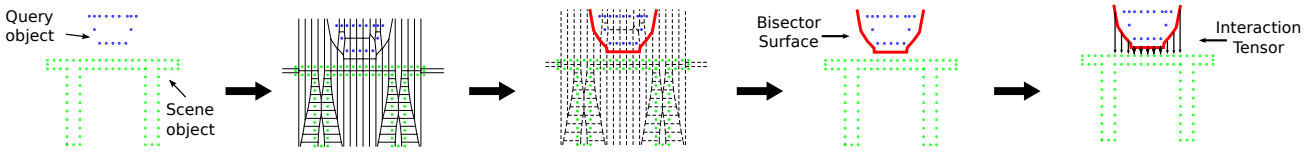

Figure 3: The interaction tensor is computed from the bisector surface. First, objects are placed simulating the interaction. The Voronoi diagram is calculated amongst all the data points. Only ridges splitting points from different objects are taken into account. These points comprise the bisector surface (red), which is used to compute the interaction tensor for *placing* a bowl on a table.

a faucet or tap and sink would act as scene-object. Fig. 3 illustrates how the interaction tensor is computed from the bisector surface between two sets of points. Specifically, it shows the iT for *placing* a bowl on a table in a simplified 2D scenario. Recall we are concerned in this paper only in the geometric component of the affordance, that is, in things and places that *appear* to afford the task.

In principle, the BS and iT extend towards infinity; in practice, we trim these to fit a sphere of radius equal to the diagonal of the query-object bounding box.

The interaction tensor inherits from the bisector surface the discriminative power in characterizing the relationships between sets of objects. It preserves key geometrical features while being robust to changes in the geometry of the interacting objects. Fig.4 shows examples of the interaction tensor for the same affordance: *placing* query-objects with changing geometries on a flat surface (table). Similarly, figure 5 shows interaction tensor examples generated using the same query-object (coat hanger) and scene-objects (coat racks) with varying geometries. In Fig.6, the same single example affordance tensor learned from the synthetic scene is used on a real RGBD scene where meaningful placements are proposed. These figures demonstrate that despite geometrical changes in the interacting objects the iT retains the overall shape or geometrical features characterizing the interaction.

One single example interaction tensor is computed for every affordance considered in our research: *placing*, *hanging*, *filling*, *sitting* and *riding*. Fig. 1 shows the training example for *riding*. Fig. 2 shows the interaction examples for the other 4 affordances.

3.2. Weighted Interaction Tensor

Every point in the BS is defined by a set *provenance vectors*, we use this information to assign a weight to every location on the interaction tensor, $W = \{w_b, \forall b \in B\}$. We assume that scene-object's point clouds are dense enough, this allows us to take simply one of such vectors without losing generality. The weight related to a point in the interaction tensor is computed from the magnitude of its corresponding *provenance vector*. This weight or distance, represents how relevant every point is for the interaction taking

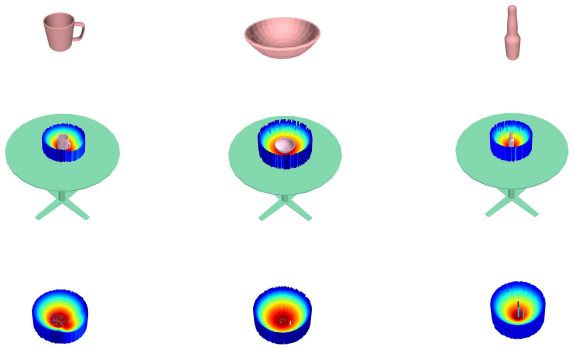

Figure 4: Examples of interaction tensor for the same affordance (*placing*) using different query-objects: mug, bowl and bottle. The interaction tensors' similarity makes them robust to changes in geometry of the query-object.

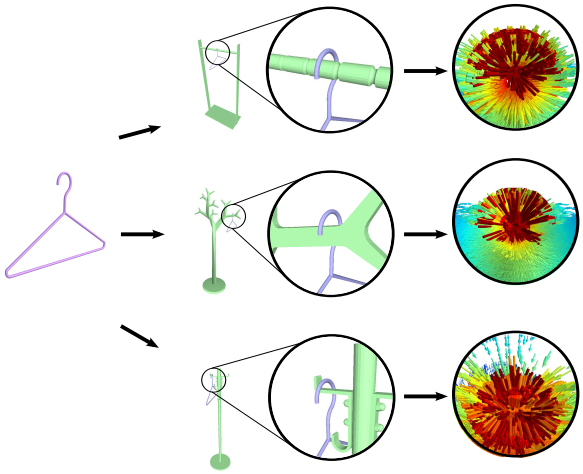

Figure 5: Interaction tensor for *hanging* a coat hanger on racks with different geometries. Although changes occur in specific locations of the tensor, the key features of the interaction are preserved.

place between the objects. Fig. 1 shows the weights for *riding* a motorcycle. Fig. 4 and 5 depict the weights for *placing* and *hanging* affordances as the color of every vector in the interaction tensor. High weights are colored in red

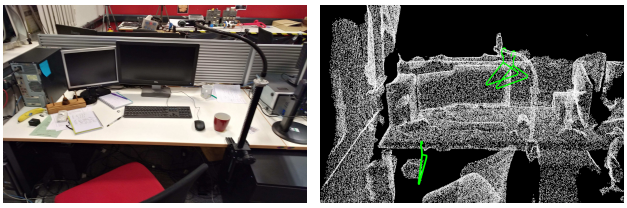

Figure 6: Affordance prediction examples for *hanging* a coat hanger in a real office-desk scene captured with an RGBD sensor.

while lower weight locations are rendered in blue.

The iT is a high dimensional and rich representation for object interactions, employing it directly as descriptor for affordance prediction would require costly computational resources. In order to reduce computational costs and improve the generalization capabilities of the descriptor, we reduce dimensionality by drawing N samples from $\mathbf{iT}$. This subset comprises what we call *affordance keypoints* $X = \{X_1, X_2, ..., X_n\}$ where $X_i = \langle b_i, \mathbf{iT}(b_i) \rangle$. This lower-dimensional descriptor is formed by a set of points on the bisector surface and **provenance vectors**. In other words, each *keypoint* is formed by a 6-dimensional feature vector which consists of the x, y, z coordinates of the data point b_i on the bisector surface, and vector $\vec{p}$ to its nearest neighbor in the scene-object $\mathbf{iT}(b_i)$. The scalar $|\vec{p}|$ encodes the importance (weight) of a keypoint in the interaction between objects; since in principle the shorter the distance between objects, the more significant the interaction is between them. Every *provenance vector* $\vec{p}$ also suggests key locations in the scene that allow the interaction to take place. Fig. 7 depicts graphically the method to compute *affordance keypoints* forming the descriptor for *placing* a bowl on a table in a 2D case.

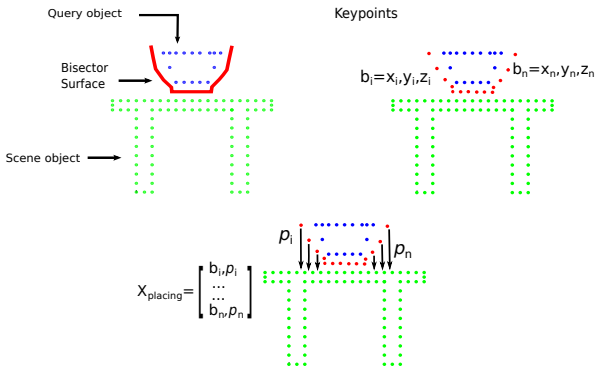

Figure 7: Affordance descriptor for *placing* a bowl on a table in a 2D scenario. A set of points is sampled from the bisector surface. An affordance keypoint is obtained by computing the interaction tensor over these sampled points. These keypoints lead to the interaction tensor descriptor X.

Each affordance descriptor X_{aff} has Nx6 dimensions, where N is the number of keypoints sampled from the iT (N=512 in our experiments). Two sampling methods of the iT were tested: 1) uniform sampling and 2) weight-driven sampling. **Weight-driven** sampling uses weights W from the tensor to form a probability distribution (2)

$$P(x_i \mid w_1, w_2, ..., w_n) = \frac{w_i}{\sum_{k=1}^{n} w_k} \tag{2}$$

$$w_i = 1 - \frac{|\vec{v_i}| - w_{min}}{w_{max} - w_{min}} \tag{3}$$

Probabilities are inversely proportional to weights after normalization and weights are given by $|\mathbf{iT}|$. Equation 3 ensures that the probabilities are in the range [0,1]. The idea behind this sampling method is to have a more meaningful representation (higher keypoint density) in locations that are highly relevant for the interaction. These typically are locations where objects come closer together or touch, for instance in Fig. 1 the saddle, handlebar grips and footrest of the motorcycle.

3.3. Affordance query

We are interested in predicting affordances or interaction possibilities on an input scene. Given a query-object and an affordance of interest, we predict good locations or possible places in the scene where the interaction could take place. Examples of such testing scenario are: "where can I place a bottle?","where can I hang a handbag? or "where can I fill a mug?"

Using these type of questions we perform a search over the input scene. Whereas this could be seen as an exhaustive process, we are able to prune the search by further characterizing each affordance with the expected orientation of the normal vector N_{aff} at the surface or point in the scene-object enabling the interaction. For instance, the *filling* affordance requires normal vectors pointing downwards in the faucet. Interactions such as *hanging* or *sitting* require normal vectors pointing upwards, which in principle come from a surface supporting the query-object.

In order to make affordance location predictions we follow Algorithm 1. First, points p_i are randomly sampled all over the input scene (30% of the total scene point in our experiments). Then, the normal vector n_i is computed at that sample point; if this normal vector is similar to the expected N_{aff} we extract a voxel centered at p_i with a radius equal to the diagonal of the query-object bounding box d_o. From the *training* example we have an approximation of the pose of X_{aff} relative to the scene-object. The transformation corresponding to that pose configuration is applied to X_{aff} to *align* it as it would be expected if the affordance could take place at p_i. Using points within the current voxel as scene-object O, a nearest-neighbor search is performed for every point b_i in X_{aff} . With this information, vectors $\vec{v_i} = \overline{b_i o_i}$

are computed at test time (*i.e.* online), these are an approximation of *provenance vectors* found in the iT example. Test vectors and example *provenance vectors* are compared to obtain a score s_i. Using n_i as rotation axis, scores are computed at different orientations θ (8 orientations evenly distributed in $[0, 2\pi)$ in our experiments). Using an empirically tuned threshold S_{aff} we are able to detect good matches (*i.e.* affordance predictions) with the most likely orientation of the query-object.

Algorithm 1 Affordance query

```
for all sample points p_i in scene do
    Compute normal vector n_i
    if n_i ≈ N_aff then
        Extract voxel of radius d_o
        for all orientations θ do
            Compute score s_i at θ_i using (4)
            if s_i ≥ S_aff then
                Predict (p_i,θ_i) as good location
```

The function to compute the *alignment* quality (*i.e.* score) at a particular location at test time is shown below

$$s_i = \sum_{i=1}^{N} \Theta(\vec{p}, \vec{v_t}) + \Delta(\vec{p}, \vec{v_t}) \tag{4}$$

where

$$\Theta(\vec{p}, \vec{v_t}) = \begin{cases} 1 & \text{if } \angle(\vec{p}\vec{v_t}) \leq W\theta_{max} \\ 0 & \text{otherwise} \end{cases}$$

and

$$\Delta(\vec{p}, \vec{v_t}) = \begin{cases} 1 & \text{if } |\vec{p}| - |\vec{v_t}| \leq W\frac{|\vec{p}|}{c} \\ 0 & \text{otherwise} \end{cases}$$

Where $W = w_{max}/w_i$, θ_{max} is the maximum angle difference allowed, c controls the maximum difference in magnitude between *provenance vectors* and test vector as a proportion of the expected distance $|\vec{p}|$. We empirically set $c = 5$ and $\theta_{max} = \frac{pi}{18}$; meaning that on more significant keypoints (highest weight) the scoring criterion in more strict and differences should not be greater than 20% of the expected values (iT example). In a first step the angle difference is computed, only if the difference between angles (Θ) is small enough the magnitudes (Δ) are compared.

4. Experimental Results

4.1. Synthetic data

For our experiments, we considered a total of fifteen synthetic scenes: 5 living rooms, 5 kitchens and 5 offices; and 8 affordance-object pairs *filling-mug, filling-cup, placing-bottle, placing-bowl, hanging-hanger, hanging-handbag, sitting-human, riding-human*. Fig. 8 show examples of the scenes that we have considered and the output affordance heat-map obtained using our algorithm. All the CAD models (objects and scenes) were publicly available from the Trimble 3D warehouse .

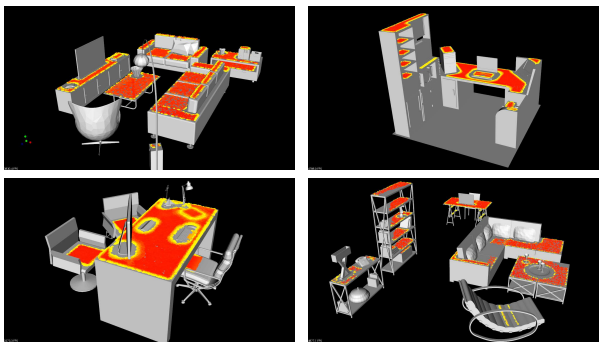

Figure 8: Example of synthetic scenes in our affordance prediction experiments. They show the prediction heat-map produced by our algorithm. The complete dataset is comprised by 5 kitchens, 5 living rooms and 5 office spaces.

Due to space limitations we only show results from a subset of our scene dataset. But more data is available upon request and in the supplementary material.

4.2. Evaluation

4.2.1 Sampling methods

In a first set of experiments we tested our sampling methods: 1) Weight-driven and 2) Uniform sampling. Fig. 9 shows the performance achieved for the 5 different affordances considered in our research. The plot shows the scores of the top 20% affordance predictions made by both methods. When weight-driven sampling is used to compute scores, the algorithm reaches the prediction score threshold faster than uniform sampling in four out of 5 affordances. This is due to the fact that by making sure that higher weight keypoints are present the algorithm is confident enough to make a prediction. High-weight keypoints are compared first by the algorithm, hence the score threshold is achieved performing less computations. On the other hand, uniform sampling achieves similar performances using on average 40% more comparisons (*i.e.* 200 keypoints more); which in most cases means all the keypoints in the descriptor. Qualitatively, both methods provide similar results, Fig 10a and 10b show this comparison for *hanging* a coat-hanger.

4.2.2 Interaction Tensor vs baselines

First we compare the performance of our approach against using the BS as descriptor. Using ICP ([25] implementa-

https://3dwarehouse.sketchup.com/

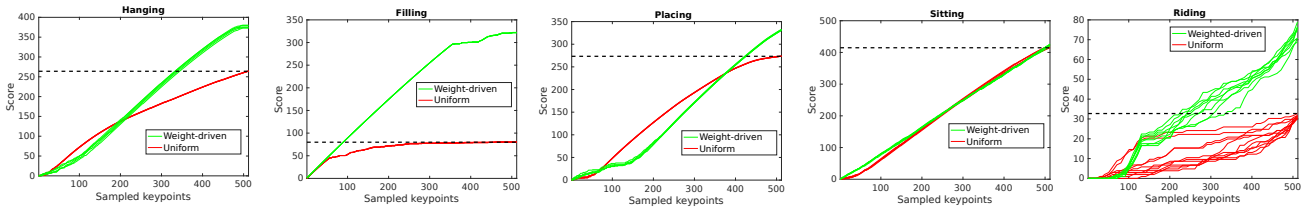

Figure 9: Plots show performance of two keypoint sampling methods. For most affordances weight-driven sampling achieves the prediction score threshold faster than uniform sampling (less comparisons made at test time). For some affordances the difference can be subtle, whereas in some others such as *filling* affordance, the difference goes to 80%.

tion), a score is computed between the BS from the interaction example and the one computed at test time. In addition to being slower or more computationally intensive, the BS descriptor is much more strict by trying to find interaction opportunities only closely similar to the *training* example. One first advantage of our approach is that, by considering a weighted vector field, we have a more relaxed matching criterion in parts of the interaction that are not critical to the affordance; this allows us to detect affordance locations in spite of variations in the scene geometry while remaining robust against false positives. In order to achieve a performance similar to the iT descriptor, it is necessary to relax the matching threshold for BS comparisons; however, this increases the number of false positives. Fig. 10d shows an example of such circumstances for *hanging* a coat-hanger on a rack.

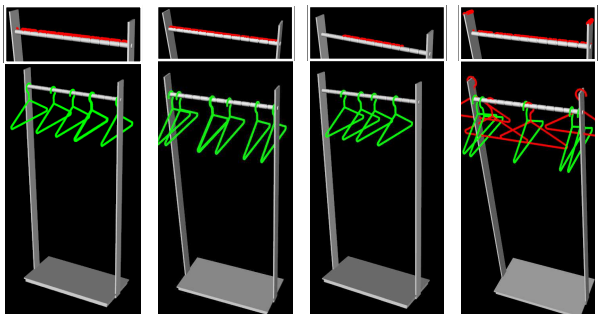

Figure 10: The iT descriptor allows more flexibility in the prediction of affordance location candidates using uniform sampling (a) and weight-driven sampling (b). The BS (c) predicts affordance location closely similar to the *training* example (center of the hanging rack). In order to achieve similar performance with BS the similarity threshold has to be relaxed (d), but at the expense of increasing the number of false positives (red coat hangers).

We then evaluated and compared our results against a baseline algorithm that we call *Naive*. This algorithm simply computes a score using ICP between the query object and the scene object, but without any explicit representation of the interaction between the objects; therefore the goal is to find the best possible alignment at test time using the score of the alignment in the interaction example as matching criteria. This is somewhat representative of methods that use object instances as examples instead of instances representing the interaction between objects. Fig. 11 shows results contrasting the *Naive* algorithm and our approach. For fairness, both of the baseline algorithms, sample points and use normal vectors on the scene similarly as we do in our approach.

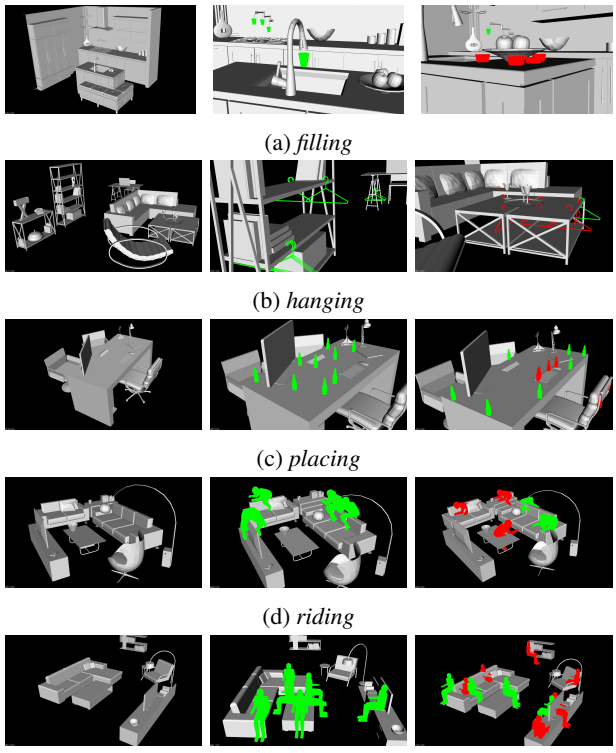

Figure 11: Affordance predictions. Results on the center column show predicted positions using the iT descriptor. Results in the column on the right show predictions made with the baseline *Naive* algorithm. *Naive* algorithm predicts good locations with equal probability as bad or unachievable configurations (red).

One of the first things to notice is that the *Naive* approach

selects good locations using mainly the normals comparison. While it does find some expected locations, it also predicts as acceptable the locations with object penetrations, occlusions or intersections; these kind of predictions would not be useful or achievable in reality. For instance 11e and 11d show *Naive* predictions for *sitting* and *riding* where the legs or parts of the body (query-object) are inside furniture. Similar cases are observed in Fig. 11a - 11c, where the predicted locations would make the query-object collide or to be inside other objects in the scene.

To further evaluate the affordance prediction results, Amazon Mechanical Turk was employed to investigate how acceptable our results are according to human criteria. Human markers were asked to select good locations for each one of the 5 affordances considered in our research. They had to choose amongst different location options, these options consisted of our top 5 and worst 5 predictions per scene, in the expectation that humans would select our top predictions over the bottom ones. A total of 85 human markers participated in the evaluation, each one provided 5 answers (one per affordance). Using this annotation as ground truth, we compute performance metrics for our top 425 predictions. Results of this evaluation are shown in Table 1, which shows that on average our approach achieves a precision of 84.92% and f-score of 91.17%; significantly outperforming the baseline methods in nearly all the affordance predictions. In other words, using a single example, our method consistently predicts top geometric affordance locations in unseen areas that agree with human criteria approximately 85 percent of the time; outperforming the baselines by 20-40% on average.

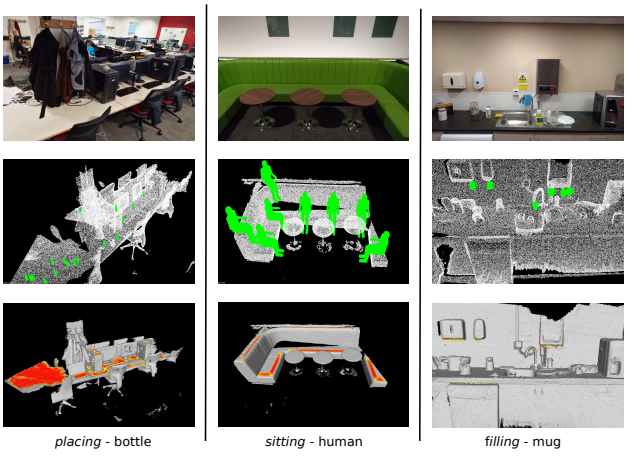

Figure 12: Affordance heatmap with predicted locations in RGBD scenes. From left to right: *placing* a bottle in office environment, *sitting* in reading room and *filling* mug in kitchen. Examples of *riding* motorbike and *hanging* coat hanger in office desk can be seen in Fig. 1 and Fig. 6 respectively.

	iT		Naive		BS	
	Accuracy	F-score	Accuracy	F-score	Accuracy	F-score
placing	96.92	**98.44**	3.08	5.97	3.08	5.97
sitting	64.62	**78.50**	64.62	**78.50**	35.38	52.27
filling	100	100	100	100	100	100
riding	92.31	**96.00**	92.31	**96.00**	7.69	14.29
hanging	70.77	**82.88**	29.23	45.24	70.77	**82.88**
Average	**84.92**	**91.17**	57.85	64.14	43.38	51.08

Table 1: Affordance prediction performance evaluated according to human markers criteria (in terms percentage).

It is worth noticing that given human location annotations for *filling*, all the algorithms predict it as good affordance locations. We believe this is mainly due to the distinctive geometry of faucets and sinks, which are usually found very seldom (one in most kitchen scenes) and this makes easier to correctly detect the *filling* affordance. In complex interactions such as *riding*, the BS algorithm is clearly outperformed by the iT descriptor. As explained previously, this algorithm mainly detects affordance at locations with scene geometries very close to the example; since there is no motorbike-like geometry it struggles to predict such affordance; a similar situations occurs for *sitting* affordances. Another remarkable result is *hanging*; according to: human judgment, iT and BS, hanging a coat hanger on edges of flat surfaces is regarded as possible. Traditional methods based on object appearance would fail to detect these cases.

4.3. RGB-D data

We conducted experiments on pointclouds captured with a Asus Xtion sensor using a publicly available dense mapping system [13]. Additionally, we included 5 publicly available scenes from [5] which contain scans of a real motorcycles and the indoor scene pointclouds from [30], leading to a testing dataset comprised by 20 real scenes. Using the same pipeline explained before, we query object-affordance pairs for each of these scenes using the *training* example from the synthetic training data. The only preprocessing step carried out to these scenes is the ground plane calibration. Fig. 12 shows affordance heat-maps for these scenes along example of the predicted locations.

5. Discussion and Conclusion

This paper presents and evaluates a new tensor field representation to express the *geometric* affordance of one object over another. By expanding the bisector surface representation to a richer tensor field, we are able to estimate affordance locations on previously unseen scenes from a single example. The introduction of weighted tensor leads to affordance keypoints that allow faster decisions per query point and a compact and straight forward way to compute a descriptor. Our evaluation is carried out with both synthetic and real RGBD scenes. The performance of our interaction tensor is significantly better in agreeing with crowdsourced opinions than the results of the baseline methods. Overall,

we see this work as an effort to motivate further advancing of approaches in Vision which, such as Active Perception [2], are more *ecological* in nature and consider the needs of the perceiving agent.

References

[1] A. Aldoma, F. Tombari, and M. Vincze. Supervised learning of hidden and non-hidden 0-order affordances and detection in real scenes. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1732–1739, May 2012. 3

[2] R. Bajcsy, Y. Aloimonos, and J. K. Tsotsos. Revisiting Active Perception. *ArXiv e-prints*, Mar. 2016. 9

[3] W. Chan, Y. Kakiuchi, K. Okada, and M. Inaba. Determining proper grasp configurations for handovers through observation of object movement patterns and inter-object interactions during usage. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, pages 1355–1360, Sept 2014. 2

[4] Y. W. Chao, Z. Wang, R. Mihalcea, and J. Deng. Mining semantic affordances of visual object categories. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4259–4267, June 2015. 2

[5] S. Choi, Q.-Y. Zhou, S. Miller, and V. Koltun. A large dataset of object scans. *arXiv:1602.02481*, 2016. 8

[6] C. Desai and D. Ramanan. Predicting functional regions on objects. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on*, pages 968–975, June 2013. 2

[7] J. J. Gibson. The theory of affordances. *Hilldale, USA*, 1977. 1, 2

[8] A. Gupta, S. Satkin, A. A. Efros, and M. Hebert. From 3d scene geometry to human workspace. In *Computer Vision and Pattern Recognition(CVPR)*, 2011. 3

[9] L. Hinkle and E. Olson. Predicting object functionality using physical simulations. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 2784–2790, Nov 2013. 3

[10] Y. Jiang and A. Saxena. Modeling high-dimensional humans for activity anticipation using gaussian process latent crfs. In *Robotics: Science and Systems*, pages 1–8, 2014. 2, 3

[11] D. Kim and G. Sukhatme. Semantic labeling of 3d point clouds with object affordance for robot manipulation. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 5578–5584, May 2014. 3

[12] H. Koppula and A. Saxena. Physically grounded spatio-temporal object affordances. In D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *Computer Vision ECCV 2014*, volume 8691 of *Lecture Notes in Computer Science*, pages 831–847. Springer International Publishing, 2014. 2

[13] S. Li and A. Calway. Rgbd relocalisation using pairwise geometry and concise key point sets. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6374–6379, May 2015. 8

[14] T. Mar, V. Tikhanoff, G. Metta, and L. Natale. Self-supervised learning of grasp dependent tool affordances on the icub humanoid robot. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3200–3206, May 2015. 2

[15] D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Freeman, 1982. 1

[16] H. Min, C. Yi, R. Luo, J. Zhu, and S. Bi. Affordance research in developmental robotics: A survey. *IEEE Transactions on Cognitive and Developmental Systems*, 8(4):237–255, Dec 2016. 2

[17] A. Myers, C. L. Teo, C. Fermüller, and Y. Aloimonos. Affordance detection of tool parts from geometric features. In *ICRA*, 2015. 3

[18] A. Nguyen, D. Kanoulas, D. G. Caldwell, and N. G. Tsagarakis. Detecting object affordances with convolutional neural networks. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2765–2770, Oct 2016. 3

[19] E. Oztop, H. Imamizu, G. Cheng, and M. Kawato. A computational model of anterior intraparietal (aip) neurons. *Neurocomput.*, 69(10-12):1354–1361, June 2006. 1

[20] A. Pandey and R. Alami. Affordance graph: A framework to encode perspective taking and effort based affordances for day-to-day human-robot interaction. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 2180–2187, Nov 2013. 2

[21] M. Peternell. Geometric properties of bisector surfaces. *Graphical Models*, 62(3):202 – 236, 2000. 2

[22] A. Pieropan, C. Ek, and H. Kjellstrom. Functional object descriptors for human activity modeling. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 1282–1289, May 2013. 2

[23] L. Piyathilaka and S. Kodagoda. Affordance-map: Mapping human context in 3d scenes using cost-sensitive svm and virtual human models. In *2015 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 2035–2040, Dec 2015. 3

[24] A. Roy and S. Todorovic. *A Multi-scale CNN for Affordance Segmentation in RGB Images*, pages 186–201. Springer International Publishing, Cham, 2016. 3

[25] R. B. Rusu and S. Cousins. 3D is here: Point Cloud Library (PCL). In *IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 9-13 2011. 6

[26] G. Saponaro, G. Salvi, and A. Bernardino. Robot anticipation of human intentions through continuous gesture recognition. In *Collaboration Technologies and Systems (CTS), 2013 International Conference on*, pages 218–225, May 2013. 2

[27] A. Srikantha and J. Gall. Discovering object classes from activities. In D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *Computer Vision – ECCV 2014*, Lecture Notes in Computer Science. Springer International Publishing, Sept. 2014. (to appear). 2

[28] W. Warren. Does this computational theory solve the right problem? marr, gibson, and the goal of vision. *Perception*, 41(9):1053–1060, 2012. 1

[29] X. Zhao, H. Wang, and T. Komura. Indexing 3d scenes using the interaction bisector surface. *ACM Trans. Graph.*, 33(3):22:1–22:14, June 2014. 2, 3

[30] Q.-Y. Zhou and V. Koltun. Dense scene reconstruction with points of interest. *ACM Trans. Graph.*, 32(4):112:1–112:8, July 2013. 8

[31] Y. Zhu, A. Fathi, and L. Fei-Fei. Reasoning about Object Affordances in a Knowledge Base Representation. In D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *Computer Vision ECCV 2014*, volume 8690 of *Lecture Notes in Computer Science*, pages 408–424. Springer International Publishing, 2014. 2